RETROSPECTIVE ANALYSIS OF COMMUNICATION EVENTS - Understanding the Dynamics of Collaborative Multi-Party Discourse.

Andrew J. Cowell

Jereme Haack

Adrienne Andrew

Rich Interaction Environments
Pacific Northwest National Laboratory
Richland, WA 99336
{andrew.cowell | jereme.haack | adrienne.andrew}@pnl.gov

Abstract

This research is aimed at understanding the dynamics of collaborative multi-party discourse across multiple communication modalities. Before we can truly make significant strides in devising collaborative communication systems, there is a need to understand how typical users utilize computationally supported communications mechanisms such as email, instant messaging, video conferencing, chat rooms, etc., both singularly and in conjunction with traditional means of communication such as face-to-face meetings, telephone calls and postal mail. Attempting to understand an individual's communications profile with access to only a single modality is challenging at best and often futile. Here, we discuss the development of RACE – Retrospective Analysis of Communications Events – a test-bed prototype to investigate issues relating to multi-modal multi-party discourse. We also examine future avenues of research for further enhancing our prototype and investigating this area.

1 Introduction

Communication is the heart of what makes us social creatures. Today, we have a myriad of technologies that allow us to communicate in ways our forefathers could never have imagined. Computationally supported modalities such as email and instant messaging have had immeasurable effect on the way we work, play and generally interact with

those in our lives. Being able to understand how individuals communicate, the methods they use, their personal preferences etc., and are all part of a field called anthroposemiotics¹. This field looks to uncover the mystery behind how we communicate with ourselves (intrapersonal communication), with others (interpersonal communication), within groups (group dynamics) and across cultures (cross-cultural communication). While there is a great deal of literature in these fields, there are few operational applications that allow for true hands-on investigation.

Perhaps nowhere is the application of this field more important than in the field of intelligence analysis. Intelligence analysts must make sound judgments, coherently constructed from scattered heterogeneous fragments of information while being faced with significant time constraints. The information they use is rarely complete, often unreliable and usually temporally and spatially diverse. These dimensions need to be aligned and the information understood to enable the analyst to recognize sequences of inter-related events and hypothesize about future actions.

Our aim has been to aid the analyst by researching, designing and implementing a test-bed for the investigation of collaborative, multi-party discourse. The focus is on reducing the complexity of analyzing communications data through a triage process; from a large corpus to a small handful of relevant conversations to finally a highly detailed view of one or two conversations, with incorporated socio-behavioral dimensions. Below we present our design methodology and discuss the latest version of the prototype.

¹ http://en.wikipedia.org/wiki/Human_communication

2 Method

The design methodology we used included a review of the literature, followed by in-depth focus group discussions with working analysts to determine requirements. Following the group session, a participatory design process was used to gather more information from our user group, leading to a set of sketches that were used in the initial prototype implementation. From these an initial prototype was created. The test-bed is currently in its second phase of implementation that includes the integration of components developed under other auspices into the RACE environment. These include indicators of affect and social roles. Finally, we have built and collected a number of data sources that we intend to use to evaluate the system. We describe these stages in the following sections.

2.1 Prior Art

This effort began with a thorough literature review across the fields of ubiquitous computing, visualization, multi-party discourse and communications theory. A number of research systems with similarities to our goal were reviewed in order to be able to understand the landscape and determine where specific opportunities may lie. Here we discuss some of the systems (mainly research prototypes) that are available for reviewing communications data.

As both Internet communications and complex graphics capabilities have become more pervasive in modern computing, there has been much interest in visualizing conversations. Due to the ease of data capture with computationally supported communications, such communication modalities as email, chat, and forum/newsgroup threads appear to be the most researched. Several systems have represented vast, multi-threaded newsgroup or forum posts such as USENET. ‘Loom’ can represent the activity patterns of individuals relative to one another, helping to characterize individuals’ participation and roles (Donath et al, 1999). In another view, linked posts are graphed to represent threads, characterizing the newsgroup as a whole. ‘Discourse Diagrams’ describes newsgroups with semantic graphs of related concepts, and also graphs people’s connectedness to one another in social networks (Sack, 2000). ‘Conversation Thumbnails’ uses an over-

view/detail display to contextualize a user’s post in the group as a whole while it is being composed (Wattenberg and Millen, 2003). ‘PeopleGarden’ represents each individual participant as a composite of their history of posting. Having all participants represented in the same screen provides insight into the dynamics of the group as a whole across its recorded history, although there is no way to track connections between individuals or threading (Xiong and Donath, 1999).

In RACE, the topics of a multitude of conversations are explored by an analyst looking for both episodic and social information. Through an iterative filtering process, the analyst examines individual conversations. Like the newsgroup visualizations above, the goals are (in addition to a general desire to understand what is going on) to determine an individual’s social role and dynamic of the group, but the concept of “conversations” is more granular. Whereas the newsgroup visualizations may represent hundreds or even thousands of users and conversation threads, the detailed visualization in RACE’s final screen represents a single discourse with as few as two people. Thus, the systems above deal with a higher level of abstraction and do not convey information on “lurkers” who may read but not post, emotional qualities of contributions, or the temporal information present in synchronous communication. RACE has the additional goals of denoting presence, affect, and what Viegas and Donath call “negotiation of conversational synchrony” (1999).

Research on chat room conversation has produced some interesting visualizations that start to deal with these concepts. The ‘Babble’ system both facilitates and visualizes synchronous and asynchronous chat (Erickson and Laff, 2001). Users are represented as colored dots on a social proxy called a ‘cookie’. The more interactions they have with the system, whether posting or only reading, the more central they become in the visualization. With inactivity, the dots move slowly back out to the periphery of the cookie, conveying information about presence and activity level. ‘Chat Circles’ is designed for synchronous chat and creates a strong sense of location by situating participants (represented as colored circles) in a large 2D space and only allowing them to see the text posted by others positioned nearby (Viegas and Donath, 1999). The circles expand to encompass posted text and shrink when ample time to

read the utterance has passed. Even people who are idling or only listening are represented spatially so others can see them. People can position their circles to avoid the ‘noise’ of unrelated conversations (as one could do at a cocktail party) or signify whom they are addressing. Each post leaves a cumulative translucent trace, indicating how long the poster has been there and how active they have been. Thus, group dynamics such as a group conversation fragmenting into smaller ones, relative verbosity, and relative position are available for interpretation.

While each of the systems above is designed for a particular modality, RACE integrates email, instant messaging, text messaging, phone conversations and teleconferences, in-person meetings in addition to chat or newsgroup participation. The goal is to get a more holistic sense of an individual throughout their discrete conversations and communication methods. As a post-hoc analysis tool, RACE aids the analyst by adding system interpretations of affect and social dynamics to the information represented in the prior art. It should be noted that this effort violates one of Erickson’s six claims about social visualization: “Portray actions, not interpretation... users understand the context better than the system ever will” (2003). We agree in theory, but the needs of our analysts differ from those of a contributor to the conversation. Content-driven interpretations of group dynamics, affect, and social role complement full-text transcripts of the conversations, providing shortcuts to insight. Below we discuss further the requirements of our user group.

2.2 Requirements Elicitation

To ensure our research was applicable to our organization’s missions and fulfilled the requirements and expectations of our user group, we enlisted the help of four analysts to determine specific requirements. These were to be our subject matter experts (SME’s). Through interactions with our SME’s we determined that while it is important to being able to understand a single conversation in time, it is just as, if not more, important to be able to comprehend the stream of conversations that occurs over longer periods, related to the same topic. For example, it is important to be able to intercept, process, and analyze a discussion between two individuals talking about making a

homemade bomb, but it is even more important to place such a discussion within the context of the set of communications leading to an understanding of the overarching plot. Such review can provide additional information that could be invaluable to the analyst. Other requirements identified as part of these sessions included:

- The system should allow the analyst to get back to original source documents and be able to review the provenance.
- The system should allow the analyst to annotate the communication events.
- Consider the use of color for note taking and marking modalities.
- The system should allow the analyst to highlight conversation fragments (i.e., small parts of a larger conversation that are considered important).
- The system should provide basic translation mechanisms for foreign language support as well as provide some form of lexicon for terms that fall outside an analyst’s field of expertise.
- The system should be able to import and export conversation fragments using common formats. The system should allow multiple analysts to work collaboratively within the same workspace.
- The system should allow the analyst to customize the environment to their preferences.

In addition to an informal list of requirements, a great deal of brainstorming was performed during this session. Following a participatory design process, system designers worked with SME’s to put together a work process and some initial sketches of the overall system that could be fed into the implementation stage.

The process was designed so that the analyst could (Figure 1) interact with the conversation corpus available to them (potentially produced as a result of a search), viewing the conversations as dots, clustered around major topics. This view could be filtered based on time period, participants involved and communications modality used.

Figure 1: Sketch of the Corpus View.

On selecting a subset of conversations to review further (Figure 2) the analyst moves through to a second screen (the sequence view) where they can analyze the conversations in relation to when they occurred (the view is reminiscent of Microsoft Project's Gantt view).

Figure 2: Sketch of the Sequence View.

While icons and text will continue to depict the modality the conversation utilized, the focus at this level is of fusing the conversations to build a sequenced stream of communications traffic so the underlying thread or purpose can be understood. Finally (Figure 3), conversations of specific interest to the analyst can be pursued in further detail in a third screen, called the 'detail view'. Here, the full transcript is displayed and can be 'played' utterance by utterance in real time. As each utterance is reached, a text-to-speech engine speaks the words, while a number of visual representations indicate social constructs such as social roles and the dynamics between the individuals.

Figure 3: Sketch of the Detail View.

2.3 Implemented Prototype

Using a participatory design process, informed by the sketches and requirements of our analysts and the limitations of current research systems, we implemented a three-screen prototype analytical environment that allows a user to visualize a large corpus of communications events (Figure 4).

Figure 4: Analyst using the RACE Environment.

The environment can run on three screens simultaneously, be split across three panes (useful for performing analysis on large displays like wall-mounted plasma displays) or on a single screen with the use of a window manager seen in the top right of each view.

For the 'corpus view' (left hand screen, Figure 5) we customized some commercially available visualization software to present the conversation corpus, clustered by topic. Zooming in to individual items brings up metadata about that specific conversation. The different modalities may also be represented by different icons or colors,

The ‘sequence view’ (Figure 6) is where we envision the majority of an analyst’s time will be spent. It is here that they will review, in detail, a small subset of conversations that they found of interest in the corpus space. For example, in their exploration of the visualization, the analyst may find a group of discussions about a particular chemical substance. Knowing that this is relevant to a study they are performing, they simply drag a box around that subset and immediately those conversations are shown in the sequence view. Each conversation has an independent time line and can be zoomed out to show the entire conversation or zoomed in to see the individual utterances (these may also be accessed using tool-tips). The conversation titles on the left hand side of the screen can be unexpanded to show all the participants involved. Clicking on the participant opens a dialog box containing known information about that individual (including any known aliases and other names they may use online). A global timeline at the bottom of the screen shows where each conversation falls in sequence.

Figure 6: The Sequence View.

Once an important conversation is uncovered through the triage process, it can be selected for deeper investigation in the details view (Figure 7). This view can enable the analyst to see beyond the individual utterances. Utilizing other research performed at the Pacific Northwest National Laboratory, the details view enables the analyst to gain insight into an individual's opinion on the topics discussed. The transcript is color-coded to show the seven dimensions of affect (expression, power, ethics, attainment, skill, accomplishment and transactions), while a graph representation allows the analyst to compare individuals' affect against each other. In order to ingest the text in different ways, a 'text-to-speech' engine can be used to have the computer 'speak' the transcript. As it steps through the utterances, a group dynamics graphic (based on Erickson's Social Proxy) shows how the individuals relate to each other, highlighting those involved in the conversation and those that are idle. This view also provides a hierarchical view of the topics discussed with the ability to trigger a multi-dimensional visualization that maps participants to topics.

Figure 7: The Detail View.

3 Evaluation and Data Sets

In addition to the prototype system, an evaluation plan was developed. The current dataset being used to demonstrate the system was synthesized from news reports about the London bombings of 7th July 2005. The evaluation will use a new dataset build up from telephone transcripts from the regional August 14, 2003 blackout² to ensure any analysts used that were involved in the development of the prototype will not benefit from any potential learning effects. This data is made up of several participants involved in many different conversations. These characteristics are exactly what RACE was designed for. Another dataset is a transcript of a murder mystery held on a chat room. While there was only space for characters to interact, there are many different threads of conversation going on at once. This data set will be useful for exploring the social dynamic part of RACE. We hope to show who the conversational “drivers” were and explore what characteristics give someone away as hiding details they do not want other characters to discover.

4 Summary & Further Work

The ultimate goal of the RACE project is to assist analysts as they try to extract meaning from a myriad of sources. To this end, we started by talking with analysts themselves. This is in recognition of the fact that no matter how powerful a tool might seem to its developers, it is useless un-

less the end users actually adopt it. By working with analysts every step of the way, we are keeping that goal in sight.

RACE’s design as a test bed enables other research to get in front of the analyst sooner. The quick insertion of the text affect work illustrates the capability to make functionality available to the user for evaluation. Showing an analyst a concrete example of an idea allows them to get a better understanding of it and an easier way to elicit feedback for future work.

While this is an exciting first step, there are many avenues of crucial research still to be performed. In many fields, having access to all the communications events that occurred is rare. Research needs to be performed to determine how best to enable the analyst to fill in these blanks. Potential approaches include hypothesized inference or the use placeholders.

Currently, the prototype analytical environment only processes and displays textual transcripts of communication events. This decision was made to handle textual content first so to ensure proof of principle prior to expending effort on the more challenging aspect of fusing video, audio, still images and text (VAST). Some effort has been expended on looking for suitable design metaphors that could aid an analyst in making sense of such diverse media (e.g., video production user interfaces such as Final Cut Pro) but more research, design and evaluation is required.

More effort needs to be expended on understanding how best to fuse different modalities of communication. Currently, a time-shifting approach is used to normalize an asynchronous email thread with similar-topic synchronous communications (e.g., telephone call, instant messaging session). This approach works but needs to be refined in order to be successful. At one level, the modality used is irrelevant – it is the essence of the event that is of primary concern. Being able to boil down the associated threads into one specific stream (e.g., multiple conversations across a number of modalities, all around the topic of plotting to explode a device at a particular location) is crucial in being able to support the analytical tradecraft and allow analysts to provide actionable intelligence to their superiors.

Conversations rarely keep to one single focused topic, and this can cause problems in the cluster visualization type approach used so far.

² <http://www.nerc.com/~filez/blackout.html>

Topic segmentation is a difficult research area and not one that we intend to pursue. There are at least three projects currently on the way at our institution that deal with this area and this work intends to utilize the fruits of those labors.

Finally, there are many elements of multi-party discourse that exist outside linguistic boundaries. The words we use, how often we make an utterance, etc., all speak to who we are as individuals. While some of this is obvious and can be observed with just a cursory review of a transcript of the source material, other elements are discrete and hidden. For example, conversational statistics can be recorded and used to determine an individual's level of engagement in a topic. Detection of familiarity (e.g., either by specific words not cur-

rently found in the present conversation or through the use of casual rather than formal speech) can indicate personal relationships between individuals in a dyad. Personality types can be inferred by markers indicative of leadership (e.g., number of interruptions performed/received, ability to change topic, use of power terms) or weaker, subversive roles (e.g., use of weak terms, submission of floor, deference to others). Analysts are rarely able to access such rich personality profiles of their subjects without performing an exhaustive analysis or calling in specialized help. While we are just beginning to integrate certain elements of social discourse, there are many other dimensions to be considered.

References

- Donath, J., Karahalios, K., and Viegas, F. 1999. "Visualizing conversations," Proceedings of HICSS-32, Persistent Conversations Track. Maui, Hawaii, January 5-8.
- Sack, W. 2000. "Discourse Diagrams: Interface Design for Very Large Scale Conversations," Proceedings of the HICSS-33, Persistent Conversations Track . Maui, Hawaii, January.
- Wattenberg, M. M & Millen, D. R., (2003) Conversation Thumbnails for Large-Scale Discussions. . In Extended Abstracts – CHI 2003. Ft. Lauderdale, April 8-10.
- Xiong, R. and Donath, J. PeopleGarden: Creating Data Portraits for Users, in Proceedings of the 12th Annual ACM Symposium on User Interface Software and Technology, New York: ACM, pp 37-44.
- Viegas, F. B. & Donath, J. S. (1999). Chat Circles. CHI 1999 Conference Proceedings: Conference on Human Factors in Computing Systems. New York: ACM Press, 9-16.
- Erickson T and Laff MR (2001) The design of the 'Babble' timeline: A social proxy for visualizing group activity over time. Extended Abstracts: Human Factors in Computing System (CHI 2001), 329-220.
- Erickson, T. Designing Visualizations of Social Activity: Six Claims, The Proceedings of CHI 2003: Extended Abstracts. New York: ACM Press, 2003.